

Electronic Structure-Based Discovery of Hybrid Photovoltaic Materials on Next-Generation HPC Platforms

Technical Report for the ALCF Theta Early Science Program

Argonne Leadership Computing Facility

ALCF Early Science Program (ESP) Technical Report

ESP Technical Reports describe the code development, porting, and optimization done in preparing an ESP project's application code(s) for the next generation ALCF computer system. This report is for a project in the Theta ESP, preparing for the ALCF Theta computer system.

About Argonne National Laboratory

Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, see www.anl.gov.

DOCUMENT AVAILABILITY

Online Access: U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free at OSTI.GOV (<http://www.osti.gov/>), a service of the U.S. Dept. of Energy's Office of Scientific and Technical Information

Reports not in digital format may be purchased by the public from the National Technical Information Service (NTIS):

U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312
www.ntis.gov
Phone: (800) 553-NTIS (6847) or (703) 605-6000
Fax: (703) 605-6900
Email: orders@ntis.gov

Reports not in digital format are available to DOE and DOE contractors from the Office of Scientific and Technical Information (OSTI):

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
www.osti.gov
Phone: (865) 576-8401
Fax: (865) 576-5728
Email: reports@osti.gov

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor UChicago Argonne, LLC, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, Argonne National Laboratory, or UChicago Argonne, LLC.

Electronic Structure-Based Discovery of Hybrid Photovoltaic Materials on Next-Generation HPC Platforms

Technical Report for the ALCF Theta Early Science Program

edited by
Timothy J. Williams and Ramesh Balakrishnan

Argonne Leadership Computing Facility

prepared by
Volker Blum, William P. Huhn, Chi (Garnett) Liu, David B. Mitzi, Yosuke Kanai, Victor Wen-zhe Yu,
Tim Rose, Noa Marom, Farren Curtis, and Álvaro Vázquez-Mayagoitia

September 2017

Electronic Structure-Based Discovery of Hybrid Photovoltaic Materials on Next-Generation HPC Platforms

Volker Blum^{1,*}, William P. Huhn¹, Chi (Garnett) Liu¹, David B. Mitzi¹, Yosuke Kanai²,
Victor Wen-zhe Yu¹, Tim Rose³, Noa Marom^{3,4}, Farren Curtis⁴, and Álvaro
Vázquez-Mayagoitia⁵

¹Duke University, Department of Mechanical Engineering and Materials Science, Durham,
NC, 27707, USA

²Department of Chemistry, University of North Carolina at Chapel Hill, Chapel Hill, NC
27599, USA

³Department of Materials Science and Engineering, Carnegie Mellon University, Pittsburgh,
PA 15213, USA

⁴Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15213, USA

⁵ Argonne Leadership Computing Facility, Argonne National Laboratory, Argonne, IL

November 19, 2018

1 Introduction

Organic and hybrid organic-inorganic materials offer significant opportunities for applications in photovoltaics (PV) and other semiconductor applications (Ref. [1]). The versatility of organic synthesis affords broad flexibility in terms of fine-tuning the properties of the molecular components. For purely organic materials, the interfaces to inorganic electrodes are a key component that determines the efficiency of eventual devices. In organic-inorganic hybrid materials, atomic-scale inorganic components are integrated with the organic functionality to form tightly integrated, crystalline semiconductor materials with potentially entirely new electronic and carrier transport characteristics. In either organic-inorganic interfaces or hybrid materials, the key properties are governed by the details of the atomic structure, which is virtually impossible to foresee based on empirical criteria. Predictive simulations of the structure and electronic properties of such materials or interfaces would thus greatly enhance our ability to identify the most promising organic-inorganic combinations faster, reducing the need for time consuming experimental synthesis attempts while navigating a complex materials space.

Key computational challenges that arise for the successful prediction of *a priori* unknown organic bulk materials, organic-inorganic interfaces, or organic-inorganic hybrid materials include: (i) efficient access to predictive potential energy surfaces with quantum-mechanical accuracy, accounting for subtle intermolecular bonding and dispersion forces, (ii) the ability to navigate large potential structure spaces in order to identify the likely low-energy structure candidates that determine the experimental properties, and (iii) predictive approaches to predict electronic excited states for

potentially large, complex structure models, employing formalisms that are often already computationally significantly more expensive than simpler quantum-mechanical approaches to the total energy applied for (i) and (ii).

The present project focuses on the optimization of two software packages for efficient, large-scale predictions of the structure and properties of organic and organic-inorganic materials and interfaces using Theta, a 9.65 petaflop supercomputing system based on Intel’s Knight’s Landing (KNL) processor architecture. The first software package is the highly accurate, scalable all-electron electronic structure package FHI-aims (Refs. [2–8]), enabling quantum-mechanical – density-functional theory (DFT) or beyond – simulations of materials with model sizes up to and above $\sim 10,000$ atoms at the time of writing. The second package is Gator, a massively parallel, first principles genetic algorithm (GA) specifically designed for structure prediction of (semi)-rigid molecular crystals. Gator is designed to fully utilize high performance computing (HPC) architectures, by spawning several parallel GA replicas that read from and write to a common population of structures. For energy evaluations and local optimization of trial structures, Gator employs dispersion-inclusive DFT by interfacing with FHI-aims (see below). In addition, large-scale simulations in FHI-aims are facilitated by the ELSI (“ELectronic Structure Infrastructure,” <http://elsi-interchange.org>) interface [9] for Kohn-Sham DFT solvers with a particular emphasis on large system sizes and current and next-generation HPC architectures.

2 Science Summary

Our scientific targets on Theta are two specific materials classes, related to organic-inorganic hybrid materials and organic materials for PV applications.

2.1 Electronic Levels in 2D Hybrid Organic-Inorganic Perovskites

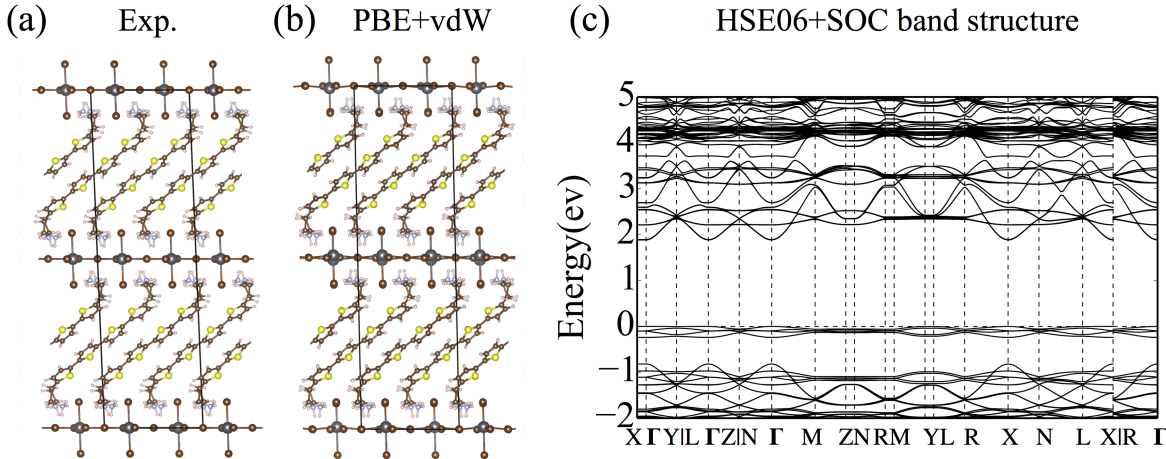


Figure 1: (a) Experimental structure of AE4TPbBr₄ (Ref. [1]). (b) PBE+vdW predicted structure of AE4TPbBr₄. (c) HSE06+SOC band structure of AE4TPbBr₄ based on the PBE+vdW predicted structure.

Organic-inorganic perovskites offer a wide range of structural flexibility (Ref. [1]). We focus particularly on so-called “2D” or “layered” perovskites, which are actually three-dimensional crystalline materials that allow one to combine complex organic functionality with a variety of inorganic components. As an example, Figures 1 (a) and (b) show the experimental and calculated structure of bis(aminoethyl)-quaterthiophene lead bromide (AE4TPbBr₄), a well-ordered crystalline material studied experimentally already in the 1990s (Ref. [10]). The full unit cell used in the simulations includes 424 atoms. The computational relaxation can be performed at the manageable semilocal DFT level (here, the Perdew-Burke-Ernzerhof (PBE) exchange-correlation functional (Ref. [11]) plus the Tkatchenko-Scheffler (TS) pairwise dispersion scheme to account for van der Waals (vdW) interactions (Ref. [12]). Electronic properties, however, require a level of theory that can qualitatively account for the fundamental gap of solids, a task that requires a much more expensive level of theory; at minimum, so-called hybrid density functionals that allow one to incorporate a fraction of the non-local exchange operator. In the example band structure shown in Fig. 1 (c), we rely on the Heyd-Scuseria-Ernzerhof (HSE06) hybrid density functional (Ref. [13, 14]) plus perturbative spin-orbit coupling (SOC) (Ref. [8]) for qualitative predictions of frontier orbital and band structure properties. For a 424-atom crystalline computational model, this is a computationally very demanding task but can be realized using the FHI-aims code on Theta. This demonstration places us in the position to explore the vast space of potential organic-inorganic hybrids with high accuracy for electronic properties, promising significant speedups in identifying hitherto unknown, tailored organic-inorganic semiconductors.

2.2 Molecular Crystal Structure Prediction Using Genetic Algorithms

Molecular crystals are a unique class of materials with diverse applications in pharmaceuticals, organic electronics, pigments, and explosives [15–24]. The molecules that comprise these crystals are bound by weak dispersion (van der Waals) interactions. As a result, the same molecule may crystallize in several different solid forms, known as polymorphs. Because the structure of a molecular crystal governs its physical properties, polymorphism may drastically impact the crystal’s desired functionality for a given application. Crystal structure prediction (CSP) is an extremely challenging task because it requires a combination of accurate electronic structure methods with efficient algorithms for configuration space exploration.

The energy differences between molecular crystal polymorphs are typically within a few kJ/mol [25–28], which calls for the accuracy of a fully quantum mechanical approach. Reaching the required accuracy has become more practical thanks to a decade of development in dispersion-inclusive DFT. In particular, the recently developed many-body dispersion (MBD) method [29–31] accurately describes the structure, energetics, dielectric properties, and mechanical properties of molecular crystals [17, 25, 32–36] by accounting for long range electrostatic screening and for non-pairwise-additive contributions of many-body dispersion interactions. For some molecular crystals the correct polymorph ranking is reproduced only when using dispersion-corrected hybrid DFT (here, the PBE0+MBD functional [25, 35]).

Genetic algorithms (GAs) can be applied to complex multidimensional search spaces, including those with many extrema or discontinuous derivatives [15], and are ideal for predicting the crystal structure of polymorphic organic molecules. They provide a good balance between exploration and exploitation by introducing an element of randomness in the mating step and subsequently locally optimizing the child structure. Furthermore, they are conceptually simple algorithms, ideal for parallelization, and can lead to unbiased and unintuitive solutions. In the context of structure

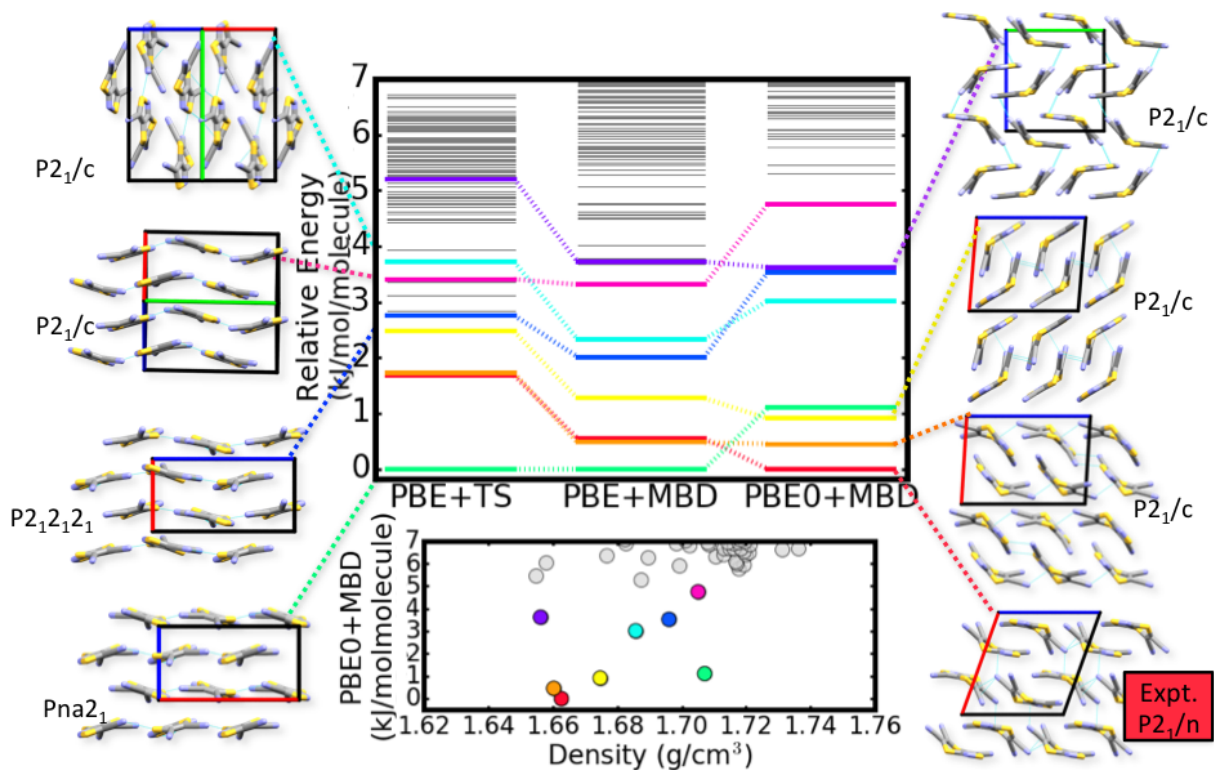


Figure 2: The top 100 crystal structures of tricyano-1,4-dithiino[c]-isothiazole as computed by Gator and reranked by various DFT functionals and dispersion corrections. Only PBE0+MBD ranks the experimental P21/n structure as the most stable crystal structure.

prediction, the target function being optimized is typically the total or free energy.

Gator successfully predicted the crystal structure of four chemically-diverse molecules selected from past CCDC blind tests: 3,4-cyclobutylfuran (Target I) (Ref. [37]), 5-cyano-3-hydroxythiophene (Target II) (Ref. [37]), 1,3-dibromo-2-chloro-5-fluorobenzene (Target XIII) (Ref. [38]), and tricyano-1,4-dithiino[c]-isothiazole (Target XXII) (Ref. [39]). The highly-accurate energetic re-ranking of the top structures for each target was performed using PBE0+MBD in FHI-aims, which is the only DFT method used that ranked the experimental structure as the number one structure for tricyano-1,4-dithiino[c]-isothiazole (Figure 5). PBE0+MBD ranked the experimental structure as first for three of the four targets studied. For the case where the experimental structure was not ranked as first, the global minimum PBE0+MBD structure is a crystal structure not previously predicted by any previous CSP study. Hybrid DFT calculations are quite computationally demanding, and given the volume of these and other DFT jobs required for successful prediction, Gator necessitated the power and speed of Theta to achieve these results.

3 Codes, Methods and Algorithms

The FHI-aims code (Refs. [2–8]) enables all-electron simulations of very high accuracy, covering materials (periodic system geometries) and molecules (zero boundary conditions at infinity) con-

taining elements across the periodic table of elements. The central numerical choice is the use of numeric atom-centered basis functions to discretize the (generalized) Kohn-Sham equations of DFT, as well as any many-body perturbation theory beyond DFT. These basis sets are flexible, capture the behavior of the orbitals near the atomic nuclei with high numerical accuracy, and are strictly localizable, lending themselves to $O(N)$ scaling implementations of grid-based integrations, density and gradient updates, and of the electrostatic potential. Several recent studies confirmed the high numerical accuracy of FHI-aims for semilocal DFT, (Refs. [8, 40–42]), hybrid DFT (Refs. [6, 42]), and many-body perturbative (Refs. [4, 40]) levels of theory. In a recent, community-wide study in the condensed matter area (Refs. [41, 43]), FHI-aims ranked on a similar level as the (linearized) augmented plane-wave (L)APW-type codes that have historically been employed as the most trusted benchmark codes in the community.

All key algorithms in FHI-aims are written to scale efficiently on massively parallel HPC architectures (Refs. [2, 3, 7, 44, 45]), relying almost exclusively on message passing interface (MPI) communication. Real-space operation such as integrals, density updates and the Hartree potential employ a non-uniform grid consisting of overlapping atom-centered grids. These operations are highly scalable thanks to a domain decomposition strategy that strictly limits the per-processor CPU time and memory consumption in the limit of very large parallel execution (Ref. [3]). The only $O(N^3)$ scaling DFT step in FHI-aims is the solution of linear generalized eigenvalue problems that arise in each self-consistency iteration. Up to several thousands of atoms, their parallel solution is facilitated by the dense eigensolver library ELPA, which provides (to the authors’ knowledge) the best available parallel scalability and execution times on distributed HPC architectures (Refs. [45, 46]). A linear-scaling approach to hybrid density functionals (Ref. [7]) based on on-the-fly screening of matrix operations and the inherent sparsity of the density matrix has previously been demonstrated to scale up to over 1,000 heavy atoms (all-electron, periodic GaAs without usage of symmetry) and up to several 1,000 CPU cores on an Intel/Infiniband type architecture. A list of HPC resources for which the authors are aware that FHI-aims has been deployed includes the current supercomputers Theta, Mira, Titan, Edison, Cory, Archer, Stampede, MareNostrum, and Sisu. The development of FHI-aims is a collaborative effort distributed throughout the world, with main centers of development located in Berlin, Durham, Munich, Helsinki, Hefei, and Graz. In the last year alone, contributions were made to the main development repo by 45 unique developers.

GATOR is written in Python with a highly modular structure that allows users to easily switch between and/or modify core GA routines for specialized purposes. GATOR offers a variety of features that enable the user to customize the search settings as needed for chemically diverse systems, including different selection, crossover, and mutation schemes. Several crossover and mutation schemes are available, acting upon structural genes, which correspond to lattice vectors, molecular positions, and molecular orientations of the parent structure(s) selected for mating.

4 Code Development

4.1 New Methods/Algorithms

Within the limited scope of the project, the effort for new methodological developments was predominantly focused on many-body perturbation theory, particularly on implementing the Bethe-Salpeter Equation (BSE) for excitonic effects in FHI-aims. This formalism allows to estimate the energy of neutral excitations based on the screened Coulomb interaction of many-body perturbation theory




Figure 3: Benchmark of FHI-aims BSE on Thiels molecular set against MOLGW.

(MBPT) in traditional solid-state physics, but had not been evaluated before in the context of numeric atom-centered basis sets. Benchmark conducted on Thiels molecular benchmark set in calculation of low-lying excitation energies gives mean absolute error (MAE) less than 1 meV compared against a recently published work (Ref. [47]) using the MOLGW code, demonstrating the accuracy of our BSE implementation (Figure 3). The convergence of predicted excitation energies was investigated for different groups of basis set: Dunning’s correlation-consistent cc-pVnZ, FHI-aims’s so-called “tier” basis sets for ground-state DFT,[2] FHI-aims’ correlation-consistent NAO-VCC-nZ basis sets for explicitly correlated methods,[48] as well as Dunning’s augmented correlation-consistent aug-cc-pVnZ basis sets. A test example for ethene, using Dunning’s aug-cc-pV5Z basis set as a reference, is shown in Figure 4. The extended augmentation functions are shown to play a key role in achieving basis converged values in non-periodic BSE calculations. Together with Jan Kloppenburg (UC Louvain, Belgium), we thus proposed a new group of basis sets, which combine the “tier 2” basis sets for DFT with two augmentation functions, demonstrating high numerical accuracy of low-lying excitations in the BSE formalism at very competitive computational cost.

A parallel implementation of the BSE was implemented within the so-called Tamm-Dancoff approximation. To improve the time and memory efficiency, future work will focus on employing the iterative Jacobi-Davidson diagonalization method to solve for low-lying excitations using the BSE matrix (as do implementations using other basis sets in the literature). These and other results will be reported in a separate, forthcoming study.

4.2 Code Refactoring

Remarkably, a scalability analysis at the outset of our FHI-aims usage on Theta demonstrated surprisingly efficient scalability of FHI-aims’ existing MPI parallelization on KNL. In particular, OpenMP optimizations that had led to speedups on the BlueGene/Q architecture Mira did not lead




Figure 4: Basis set convergence for the lowest 5 singlet excitation states of ethene calculated by BSE in FHI-aims using aug-cc-pV5Z as reference.

to faster execution times (per node) than simple single-threaded MPI execution. In our view, this observation is not entirely accidental, but is rather due to the significant amount of work invested in FHI-aims’ existing MPI parallelization in the past, which focused on memory-parallelization of essentially all significant parts, as well as a strict focus on concentrating and minimizing communication operations wherever possible. Additionally, it is the authors’ impression that the MPI library used on Theta is already very good at using shared-memory parallelization within a single node, thus further eroding a potential benefit of a simple mixed OpenMP parallelization of FHI-aims’ existing MPI parallelization.

However, operating on Theta did not always prove straightforward. In particular, the choice of the right compilers and compiler optimizations is not obvious. Using the Intel compiler and avx512-specific optimizations, a standard FHI-aims benchmark (a few MD steps for a 220-atom molecule using semilocal DFT) executed in just over 1,400 s on a single node (64 MPI tasks or ranks), compared to just over 1,300 s on a 20-core standard Intel IvyBridge node (2.5 GHz). However, for a small benchmark using hybrid DFT, the same Intel compiler and optimizations in the same binary could only complete a single evaluation of the computationally dominant exchange operator in under one hour on a single node (64 MPI tasks). On 8 nodes using 8 MPI tasks per node, this evaluation took only a few minutes – an eleven-fold speedup of the exchange operator evaluation for this particular example. While this observation is certainly related to the specific memory access pattern of the code in question, the gfortran compiler fared overall much better in our tests for hybrid DFT (see below) and was thus preferred for hybrid DFT in FHI-aims.

Within the limited scope of the project, our actual code refactoring focused on two specific areas in FHI-aims and GAtor.

4.2.1 Memory Optimizations for Calculation of Electronic Properties

We performed a code analysis of FHI-aims to identify memory intensive arrays, as proper distribution of these arrays is critical when simulating materials with 1,000s of atoms or more. One section of the code that received particular attention was the perturbative spin-orbit coupling (Ref. [8]). Here, various memory optimizations, e.g. consistent BLACS storage of eigenvectors and code restructuring to eliminate temporary arrays, were necessary to run the large 2D HOIP calculations described in Section 2.1 on Theta.

Using the information obtained from our code analysis, we implemented a light-weight memory tracking module. This infrastructure tracks the allocation and deallocation of arrays known or expected to have a significant memory footprint. The memory tracking module outputs memory usage statistics at regular intervals and at the end of the calculation. An example of the output is provided in Figure 5. This module has already seen use by developers for profiling FHI-aims’ memory footprint to suggest targeted memory optimizations, and we anticipate that the output will be useful for users to estimate the system sizes that can be achieved on computational resources.

```
Partial memory accounting:
| Residual value for overall tracked memory usage across tasks :          0.000 MB (should be 0.000 MB)
| Peak values for overall tracked memory usage (after allocating gradient_basis_wave):
|   Minimum:    1320.167 MB (on task 19)
|   Maximum:    1441.977 MB (on task 10)
|   Average:    1383.551 MB
| Largest tracked array allocation (ham_ovlp_work):
|   Minimum:    582.682 MB (on task 19)
|   Maximum:    660.525 MB (on task 10)
|   Average:    622.101 MB
Note: These values currently only include a subset of arrays which are explicitly tracked.
The "true" memory usage will be greater.
```

Figure 5: Example output for the memory tracking module implemented in FHI-aims.

A particularly memory intensive (per MPI task) array for large-scale calculations at all levels of theory (semi-local and hybrid) is the real-space Hamiltonian. The default storage format for the real-space Hamiltonian is a modification of compressed sparse row storage, known as “packed matrix” format in FHI-aims. In this simple mode of operation, the real-space integration grid is distributed across MPI tasks, but the integrated real-space Hamiltonian in the packed matrix format is synchronized across MPI tasks, so that every MPI task has a full copy of the real-space Hamiltonian. The real-space Hamiltonian will thus be the memory bottleneck for sufficiently large calculations when using the packed format.

FHI-aims therefore features an alternative distributed storage format adapted to the domain decomposition of the integration grid, a mode of operation that is (somewhat intransparently) called “local indexing” (keyword “use_local_index”). In this mode of operation, each MPI task stores only the portion of the real-space Hamiltonian to which its integration points contribute. Local indexing properly distributes the real-space Hamiltonian across MPI tasks and strictly limits the memory usage per task, removing the main memory bottleneck for calculations involving 1,000s of atoms. However, initially, the output of some interesting electronic properties obtained by post-processing after convergence of the electronic structure did not support the “local index” mode of operation, having seen its main use on architectures with more memory per core than Theta. As part of the Theta ESP, we therefore implemented local indexing support to allow us to output successfully the densities of state, Mulliken population analyses, post-processed densities of state based on denser k-space grids than the main s.c.f. cycle of FHI-aims, output of 3D arrays for eigenvector-based quantities (so-called “cube files”, covering densities, orbitals and other quantities) suitable for external visualization tools, energy band structures, and post-s.c.f. spin-orbit coupling. Memory paral-

leism was also introduced into the calculation of macroscopic dielectric functions and absorption coefficients and support for “local indexing” is planned there in future work.

4.2.2 Parallelization of GAtor

GAtor is massively parallelized. Each replica contributes to a growing, collective pool of crystal structures completely independently which enables linear scaling in the N replicas. Within each replica, new structures are generated (via breeding and mutation) and checked for uniqueness and feasibility in parallel. Then, the allotted cores on the compute nodes for each replica are partitioned to the FHI-aims binary to run the DFT calculations (colored in magenta in Figure 8).

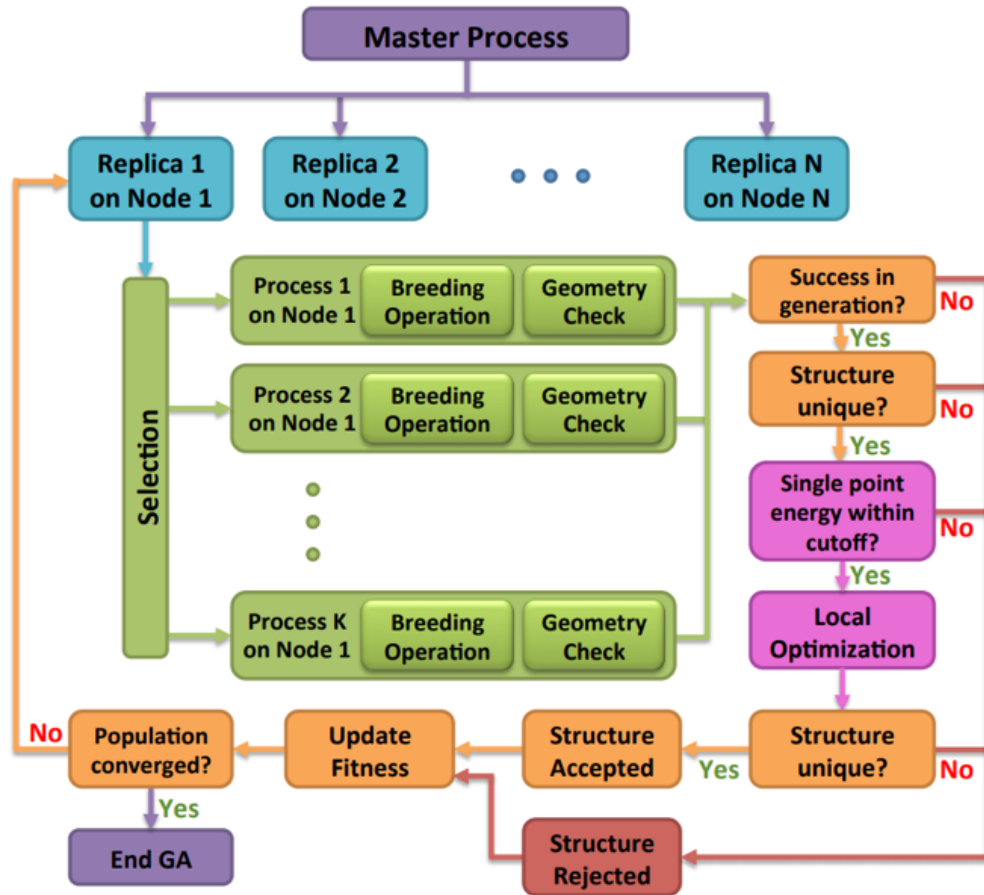


Figure 6: The parallelization scheme of GAtor.

Currently, GAtor uses the subprocess module to call FHI-aims jobs onto the already-allocated compute nodes. However, Mira’s BQ/Q CNK kernel does not support *fork* operations, binaries on the compute nodes calling other binaries. The python portions of GAtor were therefore run on Mira’s *mom* or service nodes in order to run FHI-aims on the compute nodes. However, this limited GAtor to a small number of replicas per job to prevent saturation of subprocesses on *mom* nodes and wasting some idle time on the compute nodes between calculations. The solution is to use the mpi4py module instead of the subprocess module. This allows FHI-aims to be imported into python as a library, where the MPI communicators can be passed to each of the independent FHI-aims jobs and back to the python script when the DFT calculations complete. By working

with catalysts at ALCF, we made mpi4py work on Theta and Mira’s compute nodes and can now import FHI-aims as a library. Not only does this enable Gator to truly be massively-parallel on Theta, but also obviates any runs on mom nodes and reduces compute-node idle time.

4.3 Optimization for Theta

As mentioned above, we were pleasantly surprised to be able to operate Theta efficiently for many tasks even without deep code modifications. Memory optimizations accounted for the bulk of necessary modifications to circumvent the limited amount of memory per core available on Theta. As also mentioned above, we saw specific limitations of the Intel Fortran compiler for hybrid DFT, compared to the gfortran compiler. A possible reason for this behavior is the architecture of the hybrid functional code, which relies on derived types and pointer arrays, the size of which is incremented on the fly. In contrast, the remainder of FHI-aims is deliberately written in terms of intrinsic Fortran types that are (re)allocated only infrequently, if at all, and perform much better. It is possible that the normally benign architecture of the hybrid DFT code could lead to memory fragmentation on Theta, although we did not pursue this question in depth.

The following subsections illustrate specific scalability tests performed on Theta.

4.3.1 Performance of FHI-aims on Theta




Figure 7: Scaling for FHI-aims calculations on Theta for the (a) semi-local DFT-PBE and (b) hybrid DFT-PBE0 levels of theory. 64 MPI tasks per KNL node were used, and tight integration settings and basis sets were used. DFT-PBE0 calculations were performed using executables compiled with Intel Fortran (ifort; solid lines) and gfortran (dashed lines). All DFT-PBE calculations were performed using an executable compiled with ifort. The number of atoms per computational cell are listed in parentheses in the legend of subfigure (a).

A strong-scaling benchmark of FHI-aims on Theta is shown in Figure 7 for the (a) semi-local DFT-PBE (Ref. [11]) and (b) hybrid DFT-PBEh (Ref. [49]) levels of theory. Systems sizes ranging

from 20 atoms to 1,741 atoms per unit cell are included in this benchmark. We find near-ideal scaling up to 4,096 MPI tasks at the semi-local level of theory for all systems. The ideality of scaling may be attributed to a combination of the highly vectorized batch partitioning scheme used in FHI-aims for real-space operations (Refs. [2, 3]) and the use of the dense symmetric ELPA eigensolver (Refs. [45, 46]). The reduction in scaling exponent observed for 8,192 and 16,384 MPI tasks is attributable to the ELPA eigensolver no longer scaling for sufficiently many MPI tasks, as it becomes communication-limited as the workload per processor becomes too small compared to its (necessary) fine-grained communication operations.

Calculations at the hybrid level of DFT also exhibited scaling (though sub-ideal) across nodes, consistent with the results of Ref. [7]. As noted above, We find that FHI-aims executables compiled with Intel’s ifort compiler (solid lines in subfigure(b)) exhibited notably worse scaling than executables compiled with gfortran compilers under otherwise identical conditions, although for many cases the difference is less extreme than for the specific, small test case mentioned earlier. We attribute the observed discrepancy between ifort and gfortran to an inability for ifort on Theta to properly optimize the `evaluate_exchange_matr_realspace_p0` subroutine in FHI-aims. This subroutine is specific to hybrid-functional calculations; e.g., it is not invoked for calculations in subfigure (a) which show near-ideal scaling, but it dominates the run-time for hybrid-functional calculations. It was observed that optimization flags for ifort negligibly affect the runtime for this subroutine, in stark contrast to all other HPC resources where FHI-aims has been compiled. We have not encountered this optimization issue on any other HPC resource using ifort.

4.3.2 Benchmark of AVX-512 ELPA Kernel




Figure 8: Scaling of ELPA with the AVX-512 kernel on Theta as a function of matrix dimension.

ELPA direct eigensolvers offer a massively parallel solution to symmetric/Hermitian eigenvalue for real or complex problems with improved speed and scalability to the traditional ScaLAPACK solvers, by implementing an efficient two-stage tridiagonalization algorithm and other fine-grained

numerical optimizations. To assess the performance of the ELPA solver for large-scale calculations on KNL, benchmarks were conducted with various combinations of matrix dimensions and number of KNL nodes employed on Theta. Shown in Figure 8 are timings for solving a standard eigenvalue problem with the ELPA 2-stage solver and its QR decomposition which relies in a Householder transformation kernel specifically written to use AVX-512 intrinsic functions. The parallel scalability of the ELPA solver extends to over 10^5 MPI tasks. The largest problem tested in this benchmark has a size of $1,048,576 \times 1,048,576$, which can be solved by ELPA using 3,072 of Theta’s nodes within 9,000 seconds.

4.3.3 The ELectronic Structure Infrastructure (ELSI)

The ELectronic Structure Infrastructure (ELSI) (Refs. [9, 50, 51]) provides an open-source interface to and support for several solvers to obtain the Kohn-Sham eigenvectors, eigenvalues, or the density matrix. ELSI provides a code-independent interface layer between electronic structure codes and Kohn-Sham solver libraries with automated matrix format conversion, and is now used as the default interface to solvers in FHI-aims. Theta is one of the supercomputers on which testing of ELSI was performed. Other DOE HPC resources targeted include ALCF’s Mira, OLCF’s Titan, and NERSC’s Edison and Cori.




Figure 9: Comparison of ELPA and libOMM performance on Theta. The Hamiltonian and overlap matrices used in this comparison come from a 4-layer carbon nitride model (288 atoms) calculated with FHI-aims. The five data points in each line correspond to the “tier 1”, “tier 2”, “tier 3”, “tier 4”, and “further” (beyond-tier 4) basis sets in FHI-aims, respectively.

Figure 9 shows a comparison of timings on Theta for two electronic structure solvers provided through ELSI: ELPA (version 2016.05) and the orbital minimization method libOMM (Ref. [52]). The AVX2 kernel was used for ELPA, as the AVX-512 kernel was not yet publicly available when this benchmark was performed. Results for ELPA with the AVX-512 kernel on Theta may be found

in Section 4.3.2. libOMM is an iterative method, thus the timings presented for libOMM correspond to an SCF iteration late in the SCF cycle for which the timings have stabilized.

In this benchmark, we have exploited the feature of ELSI that the various solver libraries are treated on equal footing within ELSI and can cooperate on a calculation. We compare two versions of libOMM: an older version using the ScaLAPACK pdygst subroutine and a newer version performing a similar operation with the mult_at_b_real subroutine found in ELPA. A speedup of 10x was observed in libOMM upon moving to the mult_at_b_real subroutine, demonstrating that the ability for solver libraries to cooperate with one another within ELSI allows for acceleration of calculations beyond the sum of their parts.

Installation of ELSI requires a configuration file “make.sys”, which contains installation settings for a variety of architectures and specific HPC resources. A “make.sys” file optimized for Theta is provided as part of the ELSI package, allowing users of Theta to install ELSI painlessly. Alternatively, ELSI has been integrated into and can be compiled with the current development version of FHI-aims.

5 Portability

As noted above, we did not encounter any significant portability issues of the FHI-aims code to Theta. A specific question of interest, however, is the portability to GPU based architectures. Although our Theta ESP project did not provide funding to explore this avenue in depth, we did collaborate with colleagues at Oak Ridge National Laboratory in a separate project (led by Dr. Mina Yoon) to port FHI-aims to GPUs.

Initially, exploratory work was performed into assessing the suitability of OpenARC, a general “extended OpenACC” framework, for automatically porting FHI-aims to many-core and GPU-accelerated architectures while keeping the code base unified between HPC tracks. It was determined that OpenARC would not be a suitable choice for accomplishing this objective, since significant portions of the underlying code would have had to be rewritten from Fortran 95 to C++ after all. We could also not establish speedups successfully using the OpenACC framework itself. Thus, we decided to use the conventional approach to implementation of GPU acceleration via CUDA kernels.

Two real-space operations, the electron charge density evaluation and the integration of the Hamiltonian matrix, were targeted for GPU acceleration. Shown in Figure 10 is a strong scaling benchmark for a 375-atom Bi₂Se₃ bilayer system performed on OLCF’s Titan Cray XK7. GPU-accelerated speed-ups of 1.5x to 1.6x for total timings are observed for calculations where the real-space operations dominate runtime ($\leq 2,048$ MPI tasks). The speedup is diminished for larger numbers of MPI tasks due to the end of the scaling range observed for CPU-only ELPA, which was used in this benchmark. (Note that this benchmark extends to 8,192 MPI tasks, applied for the solution of an eigenvalue problem of dimension 17,850.) Further details on our GPU work will be published in a separate manuscript.

Portability across HPC tracks was achieved by implementing a modular design, wherein all GPU acceleration is performed in special subroutines and a dummy “stub” subroutine is provided for each GPU-accelerated subroutine. HPC resources with GPU support linking against the GPU-accelerated subroutines and execute them at runtime, whereas HPC resources without GPU support link against the stub subroutines and execute the standard, e.g. non-GPU-accelerated, code path.




Figure 10: Scaling and GPU-accelerated speedups for an SCF iteration of FHI-aims for a 375-atom Bi₂Se₃ bilayer. Solid lines indicate CPU-only operation, while dashed lines of the same color as the solid lines denote the same step, but executed using GPUs. Calculations were performed on OLCF’s Titan Cray XK7. Every node utilizes 16 MPI tasks and 1 NVIDIA Tesla K20X. The presented timings are averaged over three runs. The meta-GGA TPSS functional, tight basis sets, and tight integration settings are used.

The modified stub routines cover only a few low-level areas that are numerically heavy, while the remaining code can remain entirely untouched. This strategy ensures that the FHI-aims code base remains united for GPU-accelerated and many-core HPC tracks.

6 Conclusions

Theta turns out to be a capable machine for our scientific target, that is, the electronic-structure based simulation and prediction of new complex organic and organic-inorganic hybrid materials for photovoltaic and other applications from the first principles of quantum mechanics. Remarkably, the basic code architecture of FHI-aims, built mainly upon Fortran intrinsics with relatively stable allocation patterns and a fully memory-parallel, highly optimized MPI-only architecture proved to be surprisingly effective on Theta. We note that identifying a consistent “best” set of compiler plus optimization flags did not prove trivial, as different compilers (gfortran and ifort) were found to work better for semilocal vs. hybrid DFT code parts. Thus, some care must still be taken to

evaluate the performance of Theta prior to any given run. Still, we found the ability of Theta to handle even very large and complex structure models of organic-inorganic hybrid materials at the very demanding level of hybrid DFT is impressive. Theta and computers like it may place us in the position to predict the electronic properties of semiconductor materials, their interfaces and their nanostructures correctly, with sufficient throughput, and with the full predictive power of quantum mechanics – *for the right reasons*. Computational materials science is virtually certain to see impressive advances thanks to Theta and similar machines.

Acknowledgements

This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. An award of computer time was provided by Theta Early Science program. Work on ELSI was supported by the National Science Foundation under grant number 1450280. This work was partially financially supported by the NSF through the Research Triangle MRSEC (DMR-11-21107). The GPU portion of this work was supported by the LDRD Program of ORNL managed by UT-Battelle, LLC, for the U.S. DOE and by the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. WPH and VB gratefully acknowledge the support of NVIDIA Corporation with the donation of two Tesla K40 GPGPUs and a Quadro GP100 GPGPU used for local development work.

References

- [1] Bayrammurad Saparov and David B. Mitzi. Organicinorganic perovskites: Structural versatility for functional materials design. *Chemical Reviews*, 116(7):4558–4596, 2016. doi: 10.1021/acs.chemrev.5b00715.
- [2] V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler. Ab initio molecular simulations with numeric atom-centered orbitals. *Comput. Phys. Comm.*, 180: 2175–2196, 2009. doi: 10.1016/j.cpc.2009.06.022.
- [3] V. Havu, V. Blum, P. Havu, and M. Scheffler. Efficient $O(N)$ integration for all-electron electronic structure calculation using numeric basis functions. *J. Comput. Phys.*, 228:8367–8379, 2009. doi: 10.1016/j.jcp.2009.08.008.
- [4] X. Ren, P. Rinke, V. Blum, J. Wieferink, A. Tkatchenko, A. Sanfilippo, K. Reuter, and M. Scheffler. Resolution-of-identity approach to hartreefock, hybrid density functionals, rpa, mp2 and gw with numeric atom-centered orbital basis functions. *New J. of Phys.*, 14:053020, May 2012.
- [5] Franz Knuth, Christian Carbogno, Viktor Atalla, Volker Blum, and Matthias Scheffler. All-electron formalism for total energy strain derivatives and stress tensor components for numeric atom-centered orbitals. *Computer Physics Communications*, 190(Supplement C):33 – 50, 2015. ISSN 0010-4655. doi: <https://doi.org/10.1016/j.cpc.2015.01.003>. URL <http://www.sciencedirect.com/science/article/pii/S0010465515000090>.
- [6] Arvid Conrad Ihrig, Jrgen Wieferink, Igor Ying Zhang, Matti Ropo, Xinguo Ren, Patrick Rinke, Matthias Scheffler, and Volker Blum. Accurate localized resolution of identity approach

- for linear-scaling hybrid density functionals and for many-body perturbation theory. *New J. of Phys.*, 17(9):093020, 2015. URL <http://stacks.iop.org/1367-2630/17/i=9/a=093020>.
- [7] S. V. Levchenko, X. Ren, J. Wieferink, R. Johanni, P. Rinke, V. Blum, and M. Scheffler. Hybrid functionals for large periodic systems in an all-electron, numeric atom-centered basis framework. *Comput. Phys. Comm.*, 192:60–69, 2015. doi: 10.1016/j.cpc.2015.02.021.
 - [8] W. P. Huhn and V. Blum. One-hundred-three compound band-structure benchmark of post-self-consistent spin-orbit coupling treatments in density functional theory. *Phys. Rev. Materials*, 1:033803, Aug 2017. doi: 10.1103/PhysRevMaterials.1.033803. URL <https://link.aps.org/doi/10.1103/PhysRevMaterials.1.033803>.
 - [9] V. W.-Z. Yu, F. Corsetti, A. García, W. P. Huhn, M. Jacquelin, W. Jia, B. Lange, L. Lin, J. Lu, W. Mi, A. Seifitokaldani, À. Vázquez-Mayagoitia, C. Yang, H. Yang, and V. Blum. ELSI: A unified software interface for Kohn-Sham electronic structure solvers. *Comput. Phys. Comm.*, 2017. doi: 10.1016/j.cpc.2017.09.007. URL <https://doi.org/10.1016/j.cpc.2017.09.007>. in press.
 - [10] D. B. Mitzi, K. Chondroudis, and C. R. Kagan. Design, structure, and optical properties of organic-inorganic perovskites containing an oligothiophene chromophore. *Inorganic chemistry*, 38(26):6246–6256, 1999. ISSN 1520-510X. doi: 10.1021/ic991048k. URL <http://www.ncbi.nlm.nih.gov/pubmed/11671340>.
 - [11] J. P. Perdew, K. Burke, and M. Ernzerhof. Generalized gradient approximation made simple. *Phys. Rev. Lett.*, 77:3865–3868, Oct 1996. doi: 10.1103/PhysRevLett.77.3865. URL <http://link.aps.org/doi/10.1103/PhysRevLett.77.3865>.
 - [12] T. Bučko, S. Lebègue, J. Hafner, and J. G. Ángyán. Tkatchenko-Scheffler van der Waals correction method with and without self-consistent screening applied to solids. *Phys. Rev. B*, 87:064110, Feb 2013. doi: 10.1103/PhysRevB.87.064110. URL <http://link.aps.org/doi/10.1103/PhysRevB.87.064110>.
 - [13] J. Heyd, G. E. Scuseria, and M. Ernzerhof. Hybrid functionals based on a screened Coulomb potential. *The Journal of Chemical Physics*, 118(18):8207–8215, 2003. doi: <http://dx.doi.org/10.1063/1.1564060>. URL <http://scitation.aip.org/content/aip/journal/jcp/118/18/10.1063/1.1564060>.
 - [14] J. Heyd, G. E. Scuseria, and M. Ernzerhof. Erratum: Hybrid functionals based on a screened Coulomb potential [J. Chem. Phys. 118, 8207 (2003)]. *The Journal of Chemical Physics*, 124(21):219906, 2006. doi: 10.1063/1.2204597. URL <http://dx.doi.org/10.1063/1.2204597>.
 - [15] Joel Bernstein. *Polymorphism in molecular crystals*, volume 14. Oxford University Press, 2002.
 - [16] G M Day, W D S Motherwell, and W Jones. A strategy for predicting the crystal structures of flexible molecules: the polymorphism of phenobarbital. *Phys Chem Chem Phys*, 9(14):1693–704, Apr 2007. doi: 10.1039/b612190j.
 - [17] Anthony M Reilly and Alexandre Tkatchenko. Role of dispersion interactions in the polymorphism and entropic stabilization of the aspirin crystal. *Phys Rev Lett*, 113(5):055701, Aug 2014. doi: 10.1103/PhysRevLett.113.055701.
 - [18] David P. Elder, James E. Patterson, and René Holm. The solid-state continuum: A perspective on the interrelationships between different solid-state forms in drug substance and drug

- product. *Journal of Pharmacy and Pharmacology*, 67(6):757–772, 2015. ISSN 20427158. doi: 10.1111/jphp.12293.
- [19] Colin Reese and Zhenan Bao. Organic single-crystal field-effect transistors. *Materials Today*, 10(3):20–27, 2007. ISSN 13697021. doi: 10.1016/S1369-7021(07)70016-0. URL [http://dx.doi.org/10.1016/S1369-7021\(07\)70016-0](http://dx.doi.org/10.1016/S1369-7021(07)70016-0).
 - [20] Tatsuo Hasegawa and Jun Takeya. Organic field-effect transistors using single crystals. *Science and Technology of Advanced Materials*, 10(2):024314, 2009. ISSN 1468-6996. doi: 10.1088/1468-6996/10/2/024314.
 - [21] Stefano Bergantin and Massimo Moret. Rubrene polymorphs and derivatives: The effect of chemical modification on the crystal structure. *Crystal Growth and Design*, 12(12):6035–6041, 2012. ISSN 15287483. doi: 10.1021/cg301130n.
 - [22] Pierluigi Cudazzo, Matteo Gatti, and Angel Rubio. Excitons in molecular crystals from first-principles many-body perturbation theory: Picene versus pentacene. *Phys. Rev. B*, 86:195307, Nov 2012.
 - [23] Pierluigi Cudazzo, Francesco Sottile, Angel Rubio, and Matteo Gatti. Exciton dispersion in molecular solids. *J Phys Condens Matter*, 27:113204, Mar 2015.
 - [24] Mark Fitzgerald, Michael G Gardiner, David Armit, Gregory W Dicinoski, and Craig Wall. Confirmation of the molecular structure of tetramethylene diperoxide dicarbamide (tmdd) and its sensitiveness properties. *J Phys Chem A*, 119(5):905–10, Feb 2015. doi: 10.1021/jp510827h.
 - [25] Noa Marom, Robert A. DiStasio, Viktor Atalla, Sergey Levchenko, Anthony M. Reilly, James R. Chelikowsky, Leslie Leiserowitz, and Alexandre Tkatchenko. Many-body dispersion interactions in molecular crystal polymorphism. *Angew. Chem. Int. Ed.*, 52:6629–6632, 2013.
 - [26] Aurora J Cruz-Cabeza, Susan M Reutzel-Edens, and Joel Bernstein. Facts and fictions about polymorphism. *Chemical Society Reviews*, 44(23):8619–8635, 2015.
 - [27] Gregory JO Beran. A new era for ab initio molecular crystal lattice energy prediction. *Angewandte Chemie International Edition*, 54(2):396–398, 2015.
 - [28] Gregory JO Beran. Modeling polymorphic molecular crystals with electronic structure theory. *Chem. Rev.*, 116:5567–5613, 2016.
 - [29] Robert A DiStasio, O Anatole von Lilienfeld, and Alexandre Tkatchenko. Collective many-body van der waals interactions in molecular systems. *Proceedings of the National Academy of Sciences*, 109(37):14791–14795, 2012.
 - [30] Alexandre Tkatchenko, Robert A. DiStasio, Roberto Car, and Matthias Scheffler. Accurate and efficient method for many-body van der waals interactions. *Phys. Rev. Lett.*, 108:236402, Jun 2012.
 - [31] Alberto Ambrosetti, Anthony M. Reilly, Robert A. DiStasio, and Alexandre Tkatchenko. Long-range correlation energy calculated from coupled atomic response functions. *J. Chem. Phys.*, 140:18A508, 2014.
 - [32] Bohdan Schatschneider, Jian-Jie Liang, Anthony M Reilly, Noa Marom, Guo-Xu Zhang, and Alexandre Tkatchenko. Electrodynamic response and stability of molecular crystals. *Physical Review B*, 87(6):060104, 2013.

- [33] Anthony M. Reilly and Alexandre Tkatchenko. Understanding the role of vibrations, exact exchange, and many-body van der waals interactions in the cohesive properties of molecular crystals. *J. Chem. Phys.*, 139:024705, 2013.
- [34] Alexandre Tkatchenko. Current understanding of van der waals effects in realistic materials. *Advanced Functional Materials*, 25(13):2054–2061, 2015.
- [35] Farren Curtis, Xiaopeng Wang, and Noa Marom. Effect of packing motifs on the energy ranking and electronic properties of putative crystal structures of tricyano-1,4-dithiino[c]-isothiazole. *Acta Crystallogr B Struct Sci Cryst Eng Mater*, 72(Pt 4):562–70, Aug 2016. doi: 10.1107/S2052520616009227.
- [36] Jan Hermann, Robert A DiStasio Jr, and Alexandre Tkatchenko. First-principles models for van der waals interactions in molecules and materials: Concepts, theory, and applications. 2017.
- [37] W.D.S. Motherwell, H.L. Ammon, J.D. Dunitz, A. Dzyabchenko, P. Erk, A. Gavezzottia, D.W.M. Hofmann, F.J.J. Leusen, J.P.M. Lommerse, W.T.M. Mooij, S.L. Price, H. Scheraga, B. Schweizer, M.U. Schmidt, B.P. van Eijck, P. Verwer, and D.E. Williams. Crystal structure prediction of small organic molecules: a second blind test. *Acta Crystallogr. Sect. B*, 58: 647–661, 2002.
- [38] G.M. Day, T.G. Cooper, A.J. Cruz-Cabeza, K.E. Hejczyk, H.L. Ammon, S.X.M. Boerrigter, J.S. Tan, R.G. Della Valle, E. Venuti, J. Jose, S.R. Gadre, G.R. Desiraju, T.S. Thakur, B.P. van Eijck, J.C. Facelli, V.E. Bazterra, M.B. Ferraro, D.W.M. Hofmann, M.A. Neumann, F.J.J. Leusen, J. Kendrick, S.L. Price, A.J. Misquitta, P.G. Karamertzanis, G.W.A. Welch, H.A. Scheraga, Y.A. Arnautova, M.U. Schmidt, J. van de Streek, A.K. Wolf, and B. Schweizer. Significant progress in predicting the crystal structures of small organic molecules-a report on the fourth blind test. *Acta Crystallogr. Sect. B*, 65:107–125, 2009.
- [39] A. M. Reilly, R. I. Cooper, C. S. Adjiman, S. Bhattacharya, A. D. Boese, J. G. Brandenburg, P. J. Bygrave, R. Bylsma, J. E. Campbell, R. Car, D. H. Case, R. Chadha, J. C. Cole, K. Cosburn, H. M. Cuppen, F. Curtis, G. M. Day, R. A. DiStasio Jr, A. Dzyabchenko, B. P. van Eijck, D. M. Elking, J. A. van den Ende, J. C. Facelli, M. B. Ferraro, L. Fusti-Molnar, C.-A. Gatsiou, T. S. Gee, R. de Gelder, L. M. Ghiringhelli, H. Goto, S. Grimme, R. Guo, D. W. M. Hofmann, J. Hoja, R. K. Hylton, L. Iuzzolino, W. Jankiewicz, D. T. de Jong, J. Kendrick, N. J. J. de Klerk, H.-Y. Ko, L. N. Kuleshova, X. Li, S. Lohani, F. J. J. Leusen, A. M. Lund, J. Lv, Y. Ma, N. Marom, A. E. Masunov, P. McCabe, D. P. McMahon, H. Meekes, M. P. Metz, A. J. Misquitta, S. Mohamed, B. Monserrat, R. J. Needs, M. A. Neumann, J. Nyman, S. Obata, H. Oberhofer, A. R. Oganov, A. M. Orendt, G. I. Pagola, C. C. Pantelides, C. J. Pickard, R. I. Podesszwa, L. S. Price, S. L. Price, A. Pulido, M. G. Read, K. Reuter, E. Schneider, C. Schober, G. P. Shields, P. Singh, I. J. Sugden, K. Szalewicz, C. R. Taylor, A. Tkatchenko, M. E. Tuckerman, F. Vacarro, M. Vasileiadis, . Vzquez-Mayagoitia, L. Vogt, Y. Wang, R. E. Watson, G. A. de Wijs, J. Yang, Q. Zhu, and C. R Groom. Report on the sixth blind test of organic crystal structure prediction methods. *Acta Crystallogr. Sect. B*, 72:439–459, 2016.
- [40] M. J. van Setten, F. Caruso, S. Sharifzadeh, X. Ren, M. Scheffler, F. Liu, J. Lischner, L. Lin, J. R. Deslippe, S. G. Louie, C. Yang, F. Weigend, J. B. Neaton, F. Evers, and P. Rinke. GW100: Benchmarking G0W0 for molecular systems. *J. Chem. Theory Comput.*, 11:5665–5687, 2015. doi: 10.1021/acs.jctc.5b00453.
- [41] Kurt Lejaeghere, Gustav Bihlmayer, Torbjörn Björkman, Peter Blaha, Stefan Blügel, Volker

- Blum, Damien Caliste, Ivano E Castelli, Stewart J Clark, Andrea Dal Corso, et al. Reproducibility in density functional theory calculations of solids. *Science*, 351(6280):aad3000, 2016.
- [42] S. R. Jensen, S. Saha, J. A. Flores-Livas, W. Huhn, V. Blum, S. Goedecker, and L. Frediani. The elephant in the room of density functional theory calculations. *J. Phys. Chem. Lett.*, 8: 1449–1457, 2017. doi: 10.1021/acs.jpcclett.7b00255.
 - [43] Comparing Solid State DFT Codes, Basis Sets and Potentials. URL <https://molmod.ugent.be/deltacodesdft>. <https://molmod.ugent.be/deltacodesdft>.
 - [44] L. Nemec, V. Blum, P. Rinke, and M. Scheffler. Thermodynamic equilibrium conditions of graphene films on sic. *Phys. Rev. Lett.*, 111:065502, August 2013. doi: 10.1103/PhysRevLett.111.065502.
 - [45] A. Marek, V. Blum, R. Johanni, V. Havu, T. Lang, B. Auckenthaler, A. Heinecke, H.-J. Bungartz, and H. Lederer. The ELPA library: scalable parallel eigenvalue solutions for electronic structure theory and computational science. *J. Phys.: Condens. Matter*, 26:213201, 2014. doi: 10.1088/0953-8984/26/21/213201.
 - [46] T. Auckenthaler, V. Blum, H.-J. Bungartz, T. Huckle, R. Johanni, L. Krämer, B. Lang, H. Lederer, and P. R. Willems. Parallel solution of partial symmetric eigenvalue problems from electronic structure calculations. *Parallel Computing*, 37:783–794, 2011. doi: 10.1016/j.parco.2011.05.002.
 - [47] Fabien Bruneval, Samia M. Hamed, and Jeffrey B. Neaton. A systematic benchmark of the ab initio bethe-salpeter equation approach for low-lying optical excitations of small organic molecules. *The Journal of Chemical Physics*, 142(24):244101, 2015. doi: 10.1063/1.4922489.
 - [48] Igor Ying Zhang, Xinguo Ren, Patrick Rinke, Volker Blum, and Matthias Scheffler. Numeric atom-centered-orbital basis sets with valence-correlation consistency from h to ar. *New Journal of Physics*, 15(12):123033, 2013.
 - [49] C. Adamo and V. Barone. Toward reliable density functional methods without adjustable parameters: The PBE0 model. *J. Chem. Phys.*, 110(13):6158–6170, April 1999.
 - [50] W. P. Huhn and V. W.-Z. Yu. *ELSI Interface (May 2017) User’s Guide*, 2017.
 - [51] The ELectronic Structure Infrastructure (ELSI) Interchange. URL <http://www.elsi-interchange.org>. <http://www.elsi-interchange.org>.
 - [52] F. Corsetti. The orbital minimization method for electronic structure calculations with finite-range atomic basis sets. *Comput. Phys. Comm.*, 185:873–883, 2014. doi: 10.1016/j.cpc.2013.12.008. URL <http://dx.doi.org/10.1016/j.cpc.2013.12.008>.



Argonne Leadership Computing Facility

Argonne National Laboratory
9700 South Cass Avenue, Bldg. #240
Argonne, IL 60439

www.anl.gov



Argonne National Laboratory is a U.S. Department of Energy
laboratory managed by UChicago Argonne, LLC